arXiv:0706.3794v1 [math.PR] 26 Jun 2007

On systematic scan for sampling H-colourings of the path*

Kasper Pedersen
Department of Computer Science
University of Liverpool
Liverpool L69 3BX, UK
k.pedersen@csc.liv.ac.uk

February 1, 2008

Abstract

This paper is concerned with sampling from the uniform distribution on H-colourings of the n-vertex path using systematic scan Markov chains. An H-colouring of the n-vertex path is a *homomorphism* from the n-vertex path to some fixed graph H. We show that systematic scan for H-colourings of the n-vertex path mixes in $O(\log n)$ scans for any fixed H. This is a significant improvement over the previous bound on the mixing time which was $O(n^5)$ scans. Furthermore we show that for a slightly more restricted family of H (where any two vertices are connected by a 2-edge path) systematic scan also mixes in $O(\log n)$ scans for *any scan order*. Finally, for completeness, we show that a random update Markov chain mixes in $O(n \log n)$ updates for any fixed H, improving the previous bound on the mixing time from $O(n^5)$ updates.

1 Introduction

Many combinatorial problems are of interest to computer scientists both in their own right and due to their natural applications to statistical physics. Such problems can often be studied by considering *homomorphisms* from the graph of interest G to some fixed graph H. This is known as an H-colouring of G. The vertices of H correspond to colours and the edges of H specify which colours are allowed to be adjacent in an H-colouring of a graph. Let $H = (C, E)$ by any fixed graph. We will refer to C as the set of *colours* (in the literature it is often referred to as the set of *spins*). Formally an H-colouring of a graph $G = (V, E_G)$ is a function $h : V \to C$ such that $(h(v), h(u)) \in E$ for all edges $(v, u) \in E_G$ of G. For example if H is the graph in Figure 1 then $C = \{a, b\}$ and sites (in order to be consistent with existing literature, e.g. Weitz [28], we will refer to elements of V as sites throughout this paper) assigned colour a in an H-colouring of

---

*This work was partly funded by EPSRC project GR/S76168/01.

Figure 1: The graph describing the *independent sets* model. Sites assigned colour a are "out" and sites assigned b are "in".

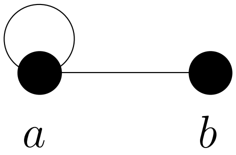

Figure 2: The graph describing the *Beach* model.

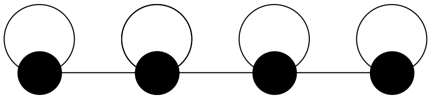

G are permitted to be adjacent to sites assigned both a and b, but sites assigned colour b can only be adjacent to sites assigned colour a.

Due to the applicability of H-colourings to models in statistical physics, and for ease of analysis, H-colouring problems are often studied by restricting attention to a specific graph H. We now give a few examples of special cases of H that correspond to important H-colouring problems. H-colourings using the graph H from Figure 1 correspond to *independent set* configurations of a graph where sites assigned colour a are "out" and sites assigned b are "in" the independent set. It is usual to assign weight 1 to vertex a and some positive weight $\lambda > 0$ to vertex b in H. Independent sets (also known as the *hard-core lattice gas model* when using the weighted setting) is one of the most commonly studied type of H-colourings in theoretical computer science. Another well-studied case is when H is the q-clique, in which case H-colourings correspond to proper q-colourings of the underlying graph. A proper q-colouring is a configuration where no two adjacent sites are permitted to be assigned the same colour. It is worth noting that proper q-colourings correspond to the q-state *anti-ferromagnetic Potts model* at zero temperature which is a well-studied model in statistical physics. Other well-known examples include the *Beach* model introduced by Burton and Steif [6] and the q-particle Widom-Rowlinson due to Widom and Rowlinson [30]. The graph corresponding to the Beach model is shown in Figure 2. The Beach model was originally introduced as an example of a physical system, with underlying graph $\mathbb{Z}^d$, which exhibits more than a single measure of maximal entropy when $d > 1$. The q-particle Widom-Rowlinson model is a model of gas consisting of q types of particles that are not allowed to be adjacent to each other. The graph corresponding to the $q = 4$ case is shown in Figure 3 where the center vertex represents empty sites and each remaining vertex represents a particle.

The problem of determining whether a graph has an H-colouring for a specific H has been well-studied and Hell and Nešetřil [20] gave a complete characterisation of graphs H for which this problem is NP-complete. In particular, they showed that if H has a loop or is bipartite then the problem is in P, and that the problem is NP-complete for any other fixed H. A complete dichotomy is also known for the problem of counting the number of H-colourings. This is due to Dyer and Greenhill [15] who showed that if H has at least one so-called *nontrivial component* then the counting problem is is #P-complete. Otherwise it is in P. A trivial component is a connected component which is either a complete graph with all loops present, or a complete bipartite graph with

Figure 3: The graph describing the 4-particle *Widom-Rowlinson* model.

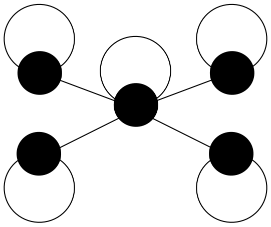

no loops present. They furthermore showed that the same dichotomy holds even when the underlying graph is of bounded degree, which is significant since in many physical applications the underlying graph tends to be of low degree. Interestingly the above characterisation for the decision problem does not hold for bounded degree graphs as shown by Galluccio, Hell and Nešetřil [17]. Despite the hardness of exactly counting the number of H-colourings it remains possible to *approximately* count the number of H-colourings of a graph as long as it is possible to sample efficiently from the (near) uniform distribution of H-colourings. This is due to a general counting-to-sampling reduction of Dyer, Goldberg and Jerrum [12] that holds for any fixed H and any underlying graph.

Sampling from the uniform distribution of H-colourings of a graph, which for this discussion we will denote by π, is a challenging task and some results about the complexity thereof are known. Goldberg, Kelk and Paterson [18] have shown that, if H has no nontrivial components, then the sampling problem is intractable in a complexity-theoretic sense. That is, they prove that there is unlikely to be any *Polynomial Almost Uniform Sampler* for H-colourings by reducing the problem of sampling from the (near) uniform distribution of H-colourings to the problem of counting independent sets in bipartite graphs, which in turn is complete for a logically-defined subclass of #P (see Dyer, Goldberg, Greenhill and Jerrum [11] for results about this complexity class). This does, however, not rule out the possibility of sampling from the uniform distribution of H-colourings of more restricted graphs, such as the n-vertex path, as we will be focusing on in this paper. This sampling task may be carried out by simulating some suitable random dynamics converging to π. Ensuring that a dynamics converges to π is generally straightforward, but obtaining good upper bounds on the number of steps required for the dynamics to become sufficiently close to π is a much more difficult problem. Due to a lack of theoretical convergence results, scientists conducting experiments by simulating such dynamics are at times forced to "guess" (using some heuristic methods) the number of steps required for their dynamics to be sufficiently close to the desired distribution. See for example Cowles and Carlin [8] for a comprehensive review of some diagnostic tools used to empirically determine these convergence rates. By establishing rigorous bounds on the convergence rates (*mixing time*) of these dynamics computer scientists can provide underpinnings for this type of experimental work and also allow a more structured approach to be taken.

Analysing the mixing time of Markov chains for H-colouring problems is a well-studied area in theoretical computer science. There is a substantial body of literature concerned with inventing Markov chains for sampling from the uniform distribution of H-colourings of graphs and providing bounds on their mixing times. When an H-colouring corresponds to a proper q-colouring of graph with maximum vertex-degree Δ then Jerrum [21], and independently Salas and Sokal [26], showed that a simple Markov chain mixes in $O(n \log n)$ updates when $q > 2\Delta$. This Markov chain makes transitions by selecting a site v and a colour c uniformly at random, and then recolouring site v to c if doing so results in a proper q-colouring of the graph. By considering a more complicated Markov chain Vigoda [27] was able to weaken the restriction on q to $q > (11/6)\Delta$ colours being sufficient for proving mixing in $O(n \log n)$ updates. This remains the least number of colours required for mixing of a Markov chain on general graphs, however the number of colours can be further reduced for special graphs. For example, when the underlying graph is the square grid then Goldberg, Martin and Paterson [19] gave a hand-proof that $q = 7$ colours are sufficient for mixing in $O(n \log n)$ updates by proving a condition called "strong spacial mixing". Achlioptas, Molloy, Moore and van Bussel [1] showed that $q = 6$ colours are sufficient for a Markov chain for proper colourings of the grid to mix in $O(n \log n)$ updates using a computational proof. As a final example for proper q-colourings Martinelli, Sinclair and Weitz [24] showed that $q = \Delta + 2$ colours are sufficient for $O(n \log n)$ mixing when the underlying graph is a tree, improving a similar result by Kenyon, Mossel and Peres [22]. When H correspond to independent set configurations of a graph with parameter λ (that is, the vertex labeled b in Figure 1 is assigned some positive weight λ and a has weight 1) then $\lambda < \frac{2}{\Delta - 2}$ is sufficient for $O(n \log n)$ mixing as shown by Dyer and Greenhill [16] and independently Luby and Vigoda [23] (although the latter result is restricted to triangle-free graphs). When $\Delta \leq 4$ these results include the $\lambda = 1$ case which is of special interest to computer scientists since it corresponds to sampling from the uniform distribution on independent sets of the graph. Weitz [29] has recently improved the condition on λ to $\lambda < (\Delta - 1)^{\Delta - 1}/(\Delta - 2)^{\Delta}$ which notably includes the $\lambda = 1$ case for $\Delta = 5$. When $\Delta \geq 6$ and $\lambda = 1$ then Dyer, Frieze and Jerrum [10] have shown that there exists a bipartite graph G_0 such that any so-called *cautious* Markov chain on independent set configurations of G_0 has (at least) exponential mixing time (in the number of sites of G_0). A Markov chain is said to be cautious if it is only allowed to change the state of a constant number of sites at the time. This negative result was generalised to H-colourings by Cooper, Dyer and Frieze [7]. Their result applies to graphs H that are either bipartite or have at least one loop present, and is not a complete graph with all loops present (observe that for such an H the decision problem is in P and the counting problem is in #P as discussed above). In particular this result guarantees the existence of a Δ-regular graph G_0 (with Δ depending on H) such that any cautious Markov chain on the set of H-colourings of G_0, and with uniform stationary distribution, has a mixing time that is at least exponential in the number of sites of G_0.

While much is now understood about the mixing times of Markov chains, the types of Markov chains frequently studied by computer scientists fall under a family of Markov chains that we call *random update Markov chains.* We say that a Markov chain on the set of H-colourings of a graph is a random update Markov chain when one step of the the process consists of randomly selecting a set of sites (often a single site) and updating the colours assigned to those sites according to some well-defined distribution induced by π. The mixing time of a random update Markov chain is measured in the number of *updates* required in order for the Markov chain to be sufficiently close to π. We point out that all the positive results described above are for random update Markov chains. An alternative to random update Markov chains is to construct a Markov chain that cycles through and updates the sites (or subsets of sites) in a deterministic order. We call this a *systematic scan Markov chain* (or *systematic scan* for short). Systematic scan may be more intuitively appealing in terms of implementation, however until recently this type of dynamics has largely resisted analysis when applied to H-colouring problems. Perhaps some of the first analyses of systematic scan were due to Amit [3] and Diaconis and Ram [9] who respectively studied systematic scan in the context of sampling from multivariate Gaussian distributions and generating random elements of a finite group. The mixing time of a systematic scan Markov chain is measured in the number of *scans* of the graph required to be sufficiently close to π and throughout this paper it holds that one scan takes $O(n)$ updates where n is the number of sites of the graph. It is important to note that systematic scan remains a random process since the method used to update the colour assigned to the selected set of sites is a randomised procedure drawing from some well-defined distribution induced by π. This paper is concerned with sampling from the uniform distribution of H-colourings of the n-vertex path using systematic scan Markov chains.

Only few results providing bounds on the mixing time of systematic scan Markov chains for sampling from the uniform distribution of H-colourings exist in the literature and almost all of them focus on proper q-colourings of bounded degree graphs. For general graphs systematic scan is known to mix in $O(\log n)$ scans whenever $q \geq 2\Delta$, where Δ is the maximum vertex-degree of the graph, by updating both end-points of an edge in each move. This is due to a recent result by Pedersen [25] which improves the polynomial in the $q = 2\Delta$ case from a result of Dyer, Goldberg and Jerrum [13] that is obtained by updating one site at the time. If the underlying graph is bipartite then a systematic scan mixes in $O(\log n)$ scans whenever $q > f(\Delta)$ where $f(\Delta) \rightarrow \beta\Delta$ as $\Delta \rightarrow \infty$ and $\beta \approx 1.76$. This result is obtained by a careful construction of the metric used in the coupling construction and is due to Bordewich, Dyer and Karpinski [4]. Furthermore, Dyer, Goldberg and Jerrum [14] have shown that a systematic scan for proper 3-colourings of the n-vertex path mixes in $\Theta(n^2 \log n)$ scans when considering a systematic scan which updates a single site at the time using the *Metropolis* update rule. In the same paper it is also shown that systematic scan for H-colourings of the n-vertex path mixes in $O(n^5)$ scans for any fixed H and that a random update Markov chain for H-colourings of the n-vertex path mixes in $O(n^5)$ updates. The authors suggest, however, that both of these bounds are unlikely to be tight and we will significantly improve them both in this paper.

In this paper we prove that systematic scan for H-colourings of the n-vertex path mixes in $O(\log n)$ scans for any fixed graph H, by updating a constant-size block of sites at each step. By constant-size we mean that the number of sites contained in a block is bounded independently of n. We do however allow the block-sizes to depend on H (since H is a fixed graph). We will present two different Markov chains in order to achieve this aim. In Section 2 we show that if H is a graph in which any pair of colours are connected by a 2-edge path then a systematic scan mixes for any order of a set of blocks, provided that the blocks are large enough. We will use a recent result by Pedersen [25], which is based on a technique known as *Dobrushin uniqueness*, in order to establish the mixing time of this Markov chain. In Section 3 we extend this result to all connected graphs H, although at the expense of imposing a specific order on the scan. The proof of mixing uses path coupling [5] in this case. Finally, for completeness, we give a proof that a random update Markov chain for H-colourings of the n-vertex path mixes in $O(n \log n)$

updates for any fixed graph H. This result is presented in Section 4.

1.1 Preliminaries and statement of results

Consider a fixed (and connected) graph $H = (C, E)$ with maximum vertex-degree Δ_H. Let $C = \{1, \ldots, q\}$ be referred to as the set of colours. Also let $V = \{1, \ldots, n\}$ be the set of sites of the n-vertex path and in particular let V_1 be the set of sites with odd indices and V_2 the set of sites with even indices. We formally say that an H-colouring of the n-vertex path is a function h from V to C such that $(h(i), h(i+1)) \in E$ for all $i \in V \setminus \{n\}$. Let Ω^+ be the set of all configurations (all possible assignments of colours to the sites) of the n-vertex path and Ω be the set of all H-colourings of the n-vertex path for the given H. Define π to be the uniform distribution on Ω. If $x \in \Omega^+$ is a configuration and $j \in V$ is a site then x_j denotes the colour assigned to j in configuration x and for any set $\Lambda \subseteq V$ let $x_\Lambda = \bigcup_{v \in \Lambda} \{x_v\}$ be the set of colours assigned to sites in Λ. For colours $c, d \in C$ and an integer l let $D^{(l)}_{c,d}$ be the uniform distribution on H-colourings of the region of consecutive sites $L = \{v_1, \ldots, v_l\} \subset V$ consistent with site v_1 being adjacent to a site $i \in V \setminus L$ assigned colour c and site v_l being adjacent to a site in $V \setminus L$ assigned colour d. Also let $D^{(l)}_{c,d}(v_j)$ be the distribution on the colour assigned to site v_j induced by $D^{(l)}_{c,d}$. Observe that for $s < l$

$$\left[D^{(l)}_{c,d} \mid v_1 = c_1, \ldots, v_s = c_s \right] = D^{(l-s)}_{c_s,d}$$

where $D^{(l)}_{c,d} \mid v_1 = c_1, \ldots, v_s = c_s$ is the uniform distribution on H-colourings of L conditioned on site v_1 being assigned colour c_1, v_2 colour c_2 and so on until v_s being assigned colour c_s.

Let $\mathcal{M}$ be any ergodic Markov chain with state space Ω and transition matrix P. By classical theory (see e.g. Aldous [2]) $\mathcal{M}$ has a unique stationary distribution, which we will denote π. The mixing time from an initial configuration $x \in \Omega$ is the number of steps, that is applications of P, required for $\mathcal{M}$ to become sufficiently close to π. Formally the mixing time of $\mathcal{M}$ from an initial configuration $x \in \Omega$ is defined, as a function of the deviation ε from stationarity, by

$$\mathrm{Mix}_x(\mathcal{M}, \varepsilon) = \min\{t > 0 \ : \ \mathrm{d_{TV}}(P^t(x, \cdot), \pi) \leq \varepsilon\},$$

where

$$\mathrm{d_{TV}}(\theta_1, \theta_2) = \frac{1}{2} \sum_i |\theta_1(i) - \theta_2(i)| = \max_{A \subseteq \Omega^+} |\theta_1(A) - \theta_2(A)|$$

is the total variation distance between two distributions θ_1 and θ_2 on Ω. The mixing time $\mathrm{Mix}(M, \varepsilon)$ of $\mathcal{M}$ is then obtained by maximising over all possible initial configurations

$$\mathrm{Mix}(\mathcal{M}, \varepsilon) = \max_{x \in \Omega} \mathrm{Mix}_x(\mathcal{M}, \varepsilon).$$

We say that $\mathcal{M}$ is *rapidly mixing* if the mixing time of $\mathcal{M}$ is polynomial in n and $\log(\varepsilon^{-1})$.

We study Markov chains that perform *heat-bath* moves on a constant number of sites at the time. For any configuration $x \in \Omega^+$ and subset of sites $\Lambda \subseteq V$ we let $\Omega_\Lambda(x)$ be the set of configurations where the colours assigned to the endpoints of each edge containing a site in Λ are also adjacent in H. A heat-bath move on Λ starting from configuration x is performed by drawing a new configuration from the uniform distribution on $\Omega_\Lambda(x)$. We would normally let Ω be the state space of our Markov chains, however, if H is bipartite then we encounter a minor technical difficulty because the Markov chain may not be ergodic. We overcome this ergodicity issue by partitioning the state space as follows. If C_1 and C_2 are the colour classes of H then $\Omega_1 = \{x \in \Omega : x_1 \in C_1\}$ is the set of H-colourings where the first site of the path is assigned a colour from C_1. Observe that in fact each site in V_1 is assigned a colour from C_1 and each site in V_2 is assigned a colour from C_2. Similarly $\Omega_2 = \{x \in \Omega : x_1 \in C_1\}$ is the set of H-colourings where the first site is assigned a colour from C_2. Intuitively, Ω_1 and Ω_2 are the two connected components of Ω and we will show (Lemma 17) that the constructed Markov chains are ergodic on either Ω_1 or Ω_2. To see that $\Omega_1 \cup \Omega_2$ contain all H-colourings of the n-vertex path it is enough to observe that if $x \in \Omega$ then any pair of adjacent sites of the n-vertex path must be assigned colours from opposite colour classes of H in x. We let $\Omega_\sim$ be the relevant state space of the Markov chains in order to ensure ergodicity. In particular, if H is non-bipartite then $\Omega_\sim = \Omega$. Otherwise H is bipartite and we let $\Omega_\sim$ be one of Ω_1 and Ω_2. This is the same partition used by Dyer et al. in [14]. See also Cooper et al. [7] for a discussion of this issue in the context of H-colourings.

We are now ready to formally define the systematic scan Markov chains we will study in this paper and state our theorems. Let $l_1 = \lceil \Delta_H^2 \log(\Delta_H^2 + 1) \rceil + 1$. Then let $\{\Theta_1, \ldots, \Theta_{m_1}\}$ be any set of $m_1 = \lceil n/l_1 \rceil$ blocks where each block consists of l_1 consecutive sites and $\bigcup_{k=1}^{m_1} \Theta_k = V$. For each block Θ_k we define $P^{[k]}$ to be the transition matrix on the state space $\Omega_\sim$ for performing a heat-bath move on Θ_k.

Definition 1. For any integer n we let $\mathcal{M}_{\text{AnyOrder}}$ be the systematic scan Markov chain, on the state space $\Omega_\sim$, with transition matrix $\Pi_{k=1}^{m_1} P^{[k]}$.

It is worth pointing out that the following result holds for *any order* of the blocks, as is the case for all results obtained by Dobrushin uniqueness (see e.g. Dyer et al [13]). In Section 2 we will use a recent result by Pedersen [25] to prove the following theorem.

Theorem 2. *Let H be a fixed connected graph and consider the systematic scan Markov chain $\mathcal{M}_{\text{AnyOrder}}$ on the state space $\Omega_\sim$. Suppose that H is a graph in which every two sites are connected by a 2-edge path. Then mixing time of $\mathcal{M}_{\text{AnyOrder}}$ is*

$$\text{Mix}(\mathcal{M}_{\text{AnyOrder}}, \varepsilon) \leq \Delta_H^2(\Delta_H^2 + 1) \log(n\varepsilon^{-1})$$

scans of the n-vertex path. This corresponds to $O(n \log n)$ updates by the construction of the set of blocks.

Remark. Note that that each H for which Theorem 2 is valid is non-bipartite so $\Omega_\sim = \Omega$.

Remark. Several well known graphs satisfy the condition of Theorem 2, for example Widom-Rowlinson configurations, independent set configurations and proper q-colourings for $q \geq 3$. The fact that an H corresponding to 3-colourings satisfies the condition of the theorem is particularly interesting since a lower bound of $\Omega(n^2 \log n)$ scans for single site systematic scan on the path is proved in Dyer at al. [14]. This means that using a simple single site coupling cannot be sufficient to establishing Theorem 2 for any family of H including 3-colourings and hence we have to use block updates.

While many natural H-colouring problems belong to the family covered by Theorem 2, others (e.g. Beach configurations) are not included. We go on to show that systematic scan mixes in $O(\log n)$ scans for any fixed graph H by placing more strict restrictions on the construction of the blocks and the order of the scan. Let $s = 4q + 1$, $\beta = \lceil \log(2sq^s + 1) \rceil q^s$ and $l_2 = 2\beta s$. For any integer n consider the following set of $m_2 + 1 = \lfloor 2n/l_2 \rfloor$ blocks $\{\Theta_0, \dots, \Theta_{m_2}\}$ where

$$\Theta_k = \{k\beta s + 1, \dots, \min((k+2)\beta s, n)\}.$$

We observe that $\bigcup_{k=0}^{m_2} \Theta_k = V$ by construction of the set of blocks. Furthermore note that the size of Θ_{m_2} is at least βs and that the size of every other block is l_2.

Definition 3. For any integer n we let $\mathcal{M}_{\text{FixedOrder}}$ be the systematic scan Markov chain, on the state space $\Omega_\sim$, which performs a heat-bath move on each block in the order $\Theta_0, \dots, \Theta_{m_2}$.

In Section 3 we will use path coupling [5] to prove the following theorem, which improves the mixing time from the corresponding result in Dyer et al. [14] from $O(n^5)$ scans to $O(\log n)$ scans.

Theorem 4. *Let H be any fixed connected graph and consider the systematic scan Markov chain $\mathcal{M}_{\text{FixedOrder}}$ on the state space $\Omega_\sim$. The mixing time of $\mathcal{M}_{\text{FixedOrder}}$ is*

$$\text{Mix}(\mathcal{M}_{\text{FixedOrder}}, \varepsilon) \leq (4sq^s + 2) \log(n\varepsilon^{-1})$$

scans of the n-vertex path. This corresponds to $O(n \log n)$ updates by the construction of the set of blocks.

Remark. It is worth remarking at this point that Theorem 4 eclipses Theorem 2 in the sense that it shows the existence of a systematic scan for a broader family of H than Theorem 2 but with the same (asymptotic) mixing time. The result stated as Theorem 2 however remains interesting in its own right since it applies to any order of the scan. Following the proof of Theorem 2 we will discuss (Observation 14) the obstacles one encounters when attempting to extend Theorem 2 to a larger family of H using the same method of proof.

For completeness we finally consider a random update Markov chain for H-colourings of the n-vertex path. Let $\gamma = 2q^s + 1$ and define the following set of $n + s\gamma - 1$ blocks, which is constructed such that each site is contained in exactly $s\gamma$ blocks

$$\Theta_k = \begin{cases} \{k, \dots, \min(k + s\gamma - 1, n)\} & \text{when } k \in \{1, \dots, n\} \\ \{1, \dots, n + s\gamma - k\} & \text{when } k \in \{n+1, \dots, n + s\gamma - 1\}. \end{cases}$$

Definition 5. For any integer n we let $\mathcal{M}_{\text{RND}}$ be the random update Markov chain, on the state space $\Omega_\sim$, which at each step selects a block uniformly at random and performs a heat-bath move on it.

In Section 4 we will use path coupling [5] to prove the following theorem, which improves the mixing time from the corresponding result in Dyer et al. [14] from $O(n^5)$ updates to $O(n \log n)$ updates.

Theorem 6. *Let H be any fixed connected graph and consider the random update Markov chain $\mathcal{M}_{\mathrm{RND}}$ on the state space $\Omega_\sim$. The mixing time of $\mathcal{M}_{\mathrm{RND}}$ is*

$$\mathrm{Mix}(\mathcal{M}_{\mathrm{RND}}, \varepsilon) \leq \frac{(n + 2sq^s + s - 1)\log(n\varepsilon^{-1})}{s}$$

block-updates. This corresponds to $O(n \log n)$ updates since the size of each block is at most $s\gamma = O(1)$.

1.2 Review of proof techniques

We now briefly introduce the techniques we will use to bound the mixing time of the above Markov chains. For technical reasons we extend the state space of the Markov chains as follows. Let Ω_1^+ be the set of configurations where each site in V_1 is assigned a colour from C_1 and each site in V_2 is assigned a colour from C_2 (recall that C_1 and C_2 are the colour classes of H). Similarly, Ω_2^+ is the set of configurations where each site in V_1 is assigned a colour from C_2 and each site in V_2 is assigned a colour from C_1. Formally

$$\Omega_1^+ = \{x \in \Omega^+ : x_{V_1} \subseteq C_1, x_{V_2} \subseteq C_2\}$$

and

$$\Omega_2^+ = \{x \in \Omega^+ : x_{V_1} \subseteq C_2, x_{V_2} \subseteq C_1\}.$$

We then extend the state space of the Markov chains to $\Omega_\sim^+$ where $\Omega_\sim^+ = \Omega^+$ if H is not bipartite and $\Omega_\sim^+$ is one of Ω_1^+ or Ω_2^+ when H is bipartite. The extended Markov chains make the same transitions as the original Markov chains on configurations in $\Omega_\sim$ and hence the extended chains do not make transitions from configurations in $\Omega_\sim$ to configurations outside $\Omega_\sim$. The stationary distributions of the extended chains are uniform over the configurations in $\Omega_\sim$ and zero elsewhere. This approach is standard and the mixing times of the original chains are bounded above by the mixing time of corresponding chain on the extended state space.

For each site $j \in V$, let S_j denote the set of pairs $(x, y) \in \Omega_\sim^+ \times \Omega_\sim^+$ of configurations that only differ on the colour assigned to site j, that is $x_i = y_i$ for all $i \neq j$. Also let $S = \bigcup_{j \in V} S_j$ be the set of all such pairs of configurations.

1.2.1 Dobrushin uniqueness

We will make use of a recent result by Pedersen [25] to prove Theorem 2 by bounding the influence *on* a site. For completeness we now summarise this result and at the same time point out how the construction of $\mathcal{M}_{\mathrm{AnyOrder}}$ ensures that all required properties are satisfied. First note from the remark after Theorem 2 that each H that we consider is not bipartite and so $\Omega_\sim^+ = \Omega^+$. Suppose that $\{\Theta_1, \dots, \Theta_m\} \subseteq V$ is a set of m blocks such that $\bigcup_{k=1}^m \Theta_k = V$ and that each block Θ_k is associated with a transition matrix $P^{[k]}$ on the state space Ω^+. For any configuration $x \in \Omega^+$, $P^{[k]}(x, \cdot)$ denotes the distribution on configurations obtained from applying $P^{[k]}$ to x. Recall from the definition of $\mathcal{M}_{\mathrm{AnyOrder}}$ that the set of blocks covers V as required and that each transition matrix $P^{[k]}$ represents performing a heat-bath move on Θ_k. It is furthermore required that each transition matrix $P^{[k]}$ satisfies the following two properties.

1. If $P^{[k]}(x, y) > 0$ then $x_i = y_i$ for each $i \in V \setminus \Theta_k$, and
2. the distribution π on Ω^+ is invariant with respect to $P^{[k]}$.

Pedersen [25] points out that if $P^{[k]}$ is the transition matrix performing a heat-bath move on Θ_k and π is the uniform distribution on Ω, as they both are in the case of $\mathcal{M}_{\text{AnyOrder}}$, then both of these properties are satisfied. These properties ensure that the stationary distribution of any systematic scan Markov chain with transition matrix $\Pi_{k=1}^m P^{[k]}$ is π.

We are now ready to define the parameter α denoting the influence on a site. For any pair of configurations $(x, y) \in S_i$ let $\Psi_k(x, y)$ be a coupling of the distributions $P^{[k]}(x, \cdot)$ and $P^{[k]}(y, \cdot)$. We remind the reader that a coupling of two distributions π_1 and π_2 on state space Ω^+ is a joint distribution on $\Omega^+ \times \Omega^+$ such that the marginal distributions are π_1 and π_2. We let $(x', y') \in \Psi_k(x, y)$ denote that the pair of configurations (x', y') is drawn from the coupling $\Psi_k(x, y)$. We then let

$$\rho_{i,j}^k = \max_{(x,y) \in S_i} \Pr_{(x',y') \in \Psi_k(x,y)}(x'_j \neq y'_j)$$

be the influence of site i on site j under Θ_k. The influence of i on j is thus the probability that site $j \in \Theta_k$ is assigned a different colour in a pair of configurations drawn from the coupling $\Psi_k(x, y)$ where x and y differ only on the colour of site i. Finally the parameter α denoting the influence on any site is defined as

$$\alpha = \max_k \max_{j \in \Theta_k} \sum_i \rho_{i,j}^k.$$

Remark. Pedersen [25] actually defines α with a positive weight assigned to each site of the graph, however as we will not use the weights in our proof they are omitted from the above definition.

The following theorem bounds the mixing time of a systematic scan Markov chain $\mathcal{M}_\rightarrow$ with transition matrix $\Pi_{k=1}^m P^{[k]}$. It is worth pointing out that, since the proof makes use of Dobrushin uniqueness, this upper-bound on the mixing time holds for any order of the blocks.

Theorem 7 (Theorem 2, Pedersen [25])**.** *If $\alpha < 1$ then the mixing time of $\mathcal{M}_\rightarrow$ satisfies*

$$\text{Mix}(\mathcal{M}_\rightarrow, \varepsilon) \leq \frac{\log(n\varepsilon^{-1})}{1 - \alpha}.$$

1.2.2 Path coupling

In order to prove Theorems 4 and 6 we will make use of path coupling [5] which is a well-known, and by now standard, technique for proving rapid mixing of Markov chains. The key idea of path coupling is to define a coupling for pairs of *adjacent* configurations where the set of all adjacent configurations connects the state space. We will say that a pair of configurations $x, y \in \Omega_\sim^+$ are adjacent if $(x, y) \in S$. The path coupling machinery then extends the coupling to all pairs of configurations in the state space. For completeness we show that S connects the state space $\Omega_\sim^+$.

Lemma 8. *The transitive closure of S is the whole of $\Omega_\sim^+ \times \Omega_\sim^+$.*

Proof. Recall that $S = \bigcup_{i \in V} S_j$ where $S_j \subseteq \Omega^+_\sim \times \Omega^+_\sim$ is the set of pairs $(x, y) \in \Omega^+_\sim \times \Omega^+_\sim$ of configurations that differ only on the colour assigned to site j. To establish the lemma it is sufficient to, for any pair of configurations $(x, y) \in \Omega^+_\sim \times \Omega^+_\sim$, to construct a path $x = z^0, z^1, \ldots, z^n = y$ such that $(z^{j-1}, z^j) \in S_j$ for each $j \in \{1, \ldots, n\}$. We define z^j for $j \in \{1, \ldots, n\}$ as follows

$$z^j_i = \begin{cases} y_i & \text{for } 1 \leq i \leq j \\ x_i & \text{for } j < i \leq n. \end{cases}$$

Informally, configuration z^j agrees with configuration y from site 1 to j and with configuration x from site $j + 1$ to n.

By definition of the configurations $z^0, \ldots, n^n$ it follows that z^{j-1} and z^j only differ on the colour assigned to site j for each $j \in \{1, \ldots, n\}$. Hence we only need to check that $z^j \in \Omega^+_\sim$ for each j. If H is non-bipartite then $\Omega^+_\sim = \Omega^+$ so $z^j \in \Omega^+_\sim$ for each $j \in \{1, \ldots, n\}$. If H is bipartite then $\Omega^+_\sim$ is one of Ω^+_1 or Ω^+_2. Suppose without loss of generality that $\Omega^+_\sim = \Omega^+_1$. Then for each $j \in \{1, \ldots n\}$ it holds by definition of Ω^+_1 that the colours x_j and y_j must be from the same colour class of H and hence have $z^j \in \Omega^+_1$. □

Finally note that $\text{Ham}(x, y) = 1$ for any $(x, y) \in S$ where $\text{Ham}(x, y)$ denotes the Hamming distance between configurations x and y. The following theorem is sufficient for our needs in this paper, and it is a special case of the general path coupling theorem proved by Bubley and Dyer [5].

Theorem 9 (Bubley, Dyer [5]). *For all pairs $(x, y) \in S$ define a coupling $(x, y) \mapsto (x', y')$ of a Markov chain $\mathcal{M}$ on the state space $\Omega^+_\sim$. Suppose that there exists a constant $0 < \gamma < 1$ such that $\mathbf{E}\left[\text{Ham}(x', y')\right] \leq (1 - \gamma)$ for all pairs $(x, y) \in S$. Then the mixing time of $\mathcal{M}$ satisfies*

$$\text{Mix}(\mathcal{M}, \varepsilon) \leq \frac{\log(n\varepsilon^{-1})}{\gamma}.$$

2 H-colourings on the path for a restricted family of H

Recall that Δ_H denotes the maximum vertex-degree of some fixed graph H and that $l_1 = \lceil \Delta^2_H \log(\Delta^2_H + 1) \rceil + 1$. The systematic scan Markov chain $\mathcal{M}_{\text{AnyOrder}}$ on $\Omega_\sim$ has transition matrix $\Pi^{m_1}_{k=1} P^{[k]}$ where $P^{[k]}$ is the transition matrix for performing a heat-bath move on block Θ_k from a set of $m_1 = \lceil n/l_1 \rceil$ size l_1 blocks covering the n-vertex path. We will prove Theorem 2, namely that $\mathcal{M}_{\text{AnyOrder}}$ mixes in $O(\log n)$ scans when H is a graph in which any two colours are connected via a 2-edge path. We will bound the mixing time of $\mathcal{M}_{\text{AnyOrder}}$ by bounding the influence on a site and begin by establishing some lemmas required to construct the coupling needed in the proof of Theorem 2.

Lemma 10. *Suppose that for any $c_1, c_2 \in C$ there is a 2-edge path in H from c_1 to c_2. Then for any $c_1, c_2, d \in C$ and integer $s' \geq 2$ there exists a coupling $\psi(D^{(s')}_{c_1,d}, D^{(s')}_{c_2,d})$ of $D^{(s')}_{c_1,d}$ and $D^{(s')}_{c_2,d}$ such that*

(i) $\Pr_{(x',y') \in \psi(D^{(s')}_{c_1,d}, D^{(s')}_{c_2,d})}(x'_{v_1} \neq y'_{v_1}) \leq 1 - \frac{1}{\Delta^2_H}$ *and*

(ii) $\Pr_{(x',y') \in \psi(D^{(2)}_{c_1,d}, D^{(2)}_{c_2,d})}(x'_{v_2} \neq y'_{v_2}) \leq 1 - \frac{1}{\Delta^2_H}$.

Proof. By the condition of the lemma there exists some $c' \in C$ adjacent to both c_1 and c_2 in H. We prove the statement by considering two cases on s.

First suppose that $s' = 2$. By the condition of the lemma there is some colour d' adjacent to both c' and d in H. There are at most Δ_H^2 valid H-colourings of the sites v_1, v_2 in either of the distributions $D^{(2)}_{c_1,d}$ and $D^{(2)}_{c_2,d}$, and hence the colouring h, which assigns c' to v_1 and d' to v_2, has weight at least $1/\Delta_H^2$ in both. We construct a coupling $\psi(D^{(2)}_{c_1,d}, D^{(2)}_{c_2,d})$ such that

$$\Pr_{(x',y')\in\psi(D^{(2)}_{c_1,d},D^{(2)}_{c_2,d})}(x' = y' = h) \geq \frac{1}{\Delta_H^2}.$$

The rest of the coupling is arbitrary. This gives the following bounds on the disagreement probabilities at v_1 and v_2

$$\Pr_{(x',y')\in\psi(D^{(2)}_{c_1,d},D^{(2)}_{c_2,d})}(x'_{v_1} = y'_{v_1}) \geq \Pr_{(x',y')\in\psi(D^{(2)}_{c_1,d},D^{(2)}_{c_2,d})}(x'_{v_1} = y'_{v_1} = c') \geq \frac{1}{\Delta_H^2}$$

which establishes (i) for $s' = 2$ and

$$\Pr_{(x',y')\in\psi(D^{(2)}_{c_1,d},D^{(2)}_{c_2,d})}(x'_{v_2} = y'_{v_2}) \geq \Pr_{(x',y')\in\psi(D^{(2)}_{c_1,d},D^{(2)}_{c_2,d})}(x'_{v_2} = y'_{v_2} = d') \geq \frac{1}{\Delta_H^2}$$

which establishes (ii).

Now suppose $s' > 2$. Let $adj(c)$ denote the set of colours adjacent to c in H and n_k the number of H-colourings on the sites $v_4, \ldots, v_{s'}$ consistent with v_3 being assigned colour $k \in C$ and $v_{s'}$ being adjacent to a site (out side the block) coloured d. Also let $p_{c,k}$ be the number of H-colourings of v_1, v_2, v_3 assigning colour c to v_1 and k to v_3 without regard to other sites. Finally let z_i be the number of H-colourings with positive measure in $D^{(s')}_{c_i,d}$ and assume without loss of generality that $z_1 \geq z_2$.

There are at most Δ_H colours available for each site in the block which gives $p_{c,k} \leq \Delta_H$ for any $c, k \in C$ and hence

$$z_1 = \sum_{c\in adj(c_1)} \sum_{k\in C} p_{c,k} n_k \leq \Delta_H \sum_{c\in adj(c_1)} \sum_{k\in C} n_k \leq \Delta_H^2 \sum_{k\in C} n_k.$$

Now let $H(c')$ be the set of all H-colourings with positive measure in $D^{(s')}_{c_1,d}$ that assign colour c' to site v_1. Let $h(c')$ denote the size of this set. Now $p_{c,k} \geq 1$ for any $c, k \in C$ since there is a 2-edge path in H between any two colours and hence

$$h(c') = \sum_{k\in C} p_{c',k} n_k \geq \sum_{k\in C} n_k.$$

Observe that, for any $h \in H(c')$, h is at least as likely in $D^{(s')}_{c_2,d}$ as in $D^{(s')}_{c_1,d}$ since we have assumed $z_1 \geq z_2$ without loss of generality. We construct a coupling $\psi(D^{(s')}_{c_1,d}, D^{(s')}_{c_2,d})$ of $D^{(s')}_{c_1,d}$ and $D^{(s')}_{c_2,d}$ in which for each $h \in H(c')$

$$\Pr_{(x',y')\in\psi(D^{(s')}_{c_1,d},D^{(s')}_{c_2,d})}(x' = y' = h) \geq \frac{1}{z_1}.$$

The rest of the coupling is arbitrary. Hence

$$\begin{aligned}\Pr_{(x',y')\in\psi(D^{(s')}_{c_1,d},D^{(s')}_{c_2,d})}(x'_{v_1}=y'_{v_1}) &\geq \sum_{h\in H(c')}\Pr_{(x',y')\in\psi(D^{(s')}_{c_1,d},D^{(s')}_{c_2,d})}(x'=y'=h)\\ &\geq \frac{h(c')}{z_1}\\ &\geq \frac{1}{\Delta_H^2}\end{aligned}$$

using the bounds on z_1 and $h(c')$. This completes the proof. ☐

We then use Lemma 10 to bound the disagreement probabilities at each site of of the block when a pair of configurations are drawn from a recursively constructed coupling.

Lemma 11. *Suppose that for any $c_1, c_2 \in C$ there is a 2-edge path in H from c_1 to c_2. Then for all $c_1, c_2, d \in C$ and integers $l' \geq 2$ there exists a coupling $\Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})$ of $D^{(l')}_{c_1,d}$ and $D^{(l')}_{c_2,d}$ in which for $j \in \{1,\dots,l'-1\}$*

$$\Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_j}\neq y'_{v_j}) \leq \left(1-\frac{1}{\Delta_H^2}\right)^j$$

and

$$\Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_l}\neq y'_{v_l}) \leq \left(1-\frac{1}{\Delta_H^2}\right)^{l'-1}.$$

Proof. We recursively construct a coupling $\Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})$ of $D^{(l')}_{c_1,d}$ and $D^{(l')}_{c_2,d}$ using the method set out in Goldberg et al. [19] as follows. Firstly $l' = 2$ is the base case and we use the coupling from Lemma 10. For $l' \geq 3$ we construct a coupling using the following two step process.

1. Couple $D^{(l')}_{c_1,d}(v_1)$ and $D^{(l')}_{c_2,d}(v_1)$ greedily to maximise the probability of assigning the same colour to site v_1 in both distributions.
2. If the same colour c was chosen for v_1 in both distributions in step 1 then the set of valid H-colourings of the remaining sites are the same in both distributions. Hence the conditional distributions $D^{(l')}_{c_1,d} \mid v_1 = c$ and $D^{(l')}_{c_2,d} \mid v_1 = c$ are the same and the rest of the coupling is trivial. Otherwise, for all pairs (c'_1, c'_2) of distinct colours recursively couple $D^{(l')}_{c_1,d} \mid v_1 = c'_1 = D^{(l'-1)}_{c'_1,d}$ and $D^{(l')}_{c_2,d} \mid v_1 = c'_2 = D^{(l'-1)}_{c'_2,d}$ which is a sub problem of size $l'-1$.

This completes the coupling construction.

Now for $j \in \{1,\dots,l'-1\}$ we prove by induction that

$$\Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_j}\neq y'_{v_j}) \leq \left(1-\frac{1}{\Delta_H^2}\right)^j. \tag{1}$$

The base case, $j = 1$, follows from Lemma 10 since we couple the colour at site v_1 greedily to maximise the probability of agreement at v_1 in the first step of the recursive coupling. Now suppose that (1) is true for $j - 1$ then

$$\begin{aligned}
&\Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_j} \neq y'_{v_j}) \\
&\quad = \sum_{c'_1,c'_2} \Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_{j-1}} = c'_1, y'_{v_{j-1}} = c'_2) \\
&\qquad \times \Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d}|v_{j-1}=c'_1,D^{(l')}_{c_2,d}|v_{j-1}=c'_2)}(x'_{v_j} \neq y'_{v_j}) \\
&\quad = \sum_{c'_1,c'_2} \Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_{j-1}} = c'_1, y'_{v_{j-1}} = c'_2) \\
&\qquad \times \Pr_{(x',y')\in\Psi(D^{(l'-j+1)}_{c'_1,d},D^{(l'-j+1)}_{c'_2,d})}(x'_{v_1} \neq y'_{v_1}) \\
&\quad \leq \sum_{c'_1,c'_2} \Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_{j-1}} = c'_1, y'_{v_{j-1}} = c'_2) \left(1 - \frac{1}{\Delta_H^2}\right) \\
&\quad \leq \left(1 - \frac{1}{\Delta_H^2}\right)^j
\end{aligned}$$

where the first inequality uses Lemma 10 and the second is the inductive hypothesis.

The $j = l'$ case is similar.

$$\begin{aligned}
&\Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_l} \neq y'_{v_l}) \\
&\quad = \sum_{c'_1,c'_2} \Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_{l'-2}} = c'_1, y'_{v_{l'-2}} = c'_2) \\
&\qquad \times \Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d}|v_{l'-2}=c'_1,D^{(l')}_{c_2,d}|v_{l'-2}=c'_2)}(x'_{v_{l'}} \neq y'_{v_{l'}}) \\
&\quad = \sum_{c'_1,c'_2} \Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_{l'-2}} = c'_1 \wedge y'_{v_{l'-2}} = c'_2) \\
&\qquad \times \Pr_{(x',y')\in\Psi(D^{(2)}_{c'_1,d},D^{(2)}_{c'_2,d})}(x'_{v_2} \neq y'_{v_2}) \\
&\quad \leq \left(1 - \frac{1}{\Delta_H^2}\right)^{l'-2} \left(1 - \frac{1}{\Delta_H^2}\right) = \left(1 - \frac{1}{\Delta_H^2}\right)^{l'-1}
\end{aligned}$$

where the inequality uses Lemma 10 and (1). □

We can then use the coupling constructed in Lemma 11 to construct a coupling $\Psi_k(x, y)$ of the distributions $P^{[k]}(x, \cdot)$ and $P^{[k]}(y, \cdot)$ for each pair of configurations $(x, y) \in S_i$. We summarise the disagreement probabilities in this coupling in the following corollary (of Lemma 11).

Corollary 12. *For any sites $i, j \in V$ let $d(i, j)$ denote the edge distance between them and suppose that for any $c, d \in C$ there exists a 2-edge path in H from c to d. Then*

$$\rho^k_{i,j} \leq \begin{cases} \left(1 - \frac{1}{\Delta_H^2}\right)^{d(i,j)} & \text{if } i \text{ is on the boundary of } \Theta_k \text{ and } d(i,j) < l_1 \\ \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - 1} & \text{if } i \text{ is on the boundary of } \Theta_k \text{ and } d(i,j) = l_1 \\ 0 & \text{otherwise.} \end{cases}$$

Figure 4: A block Θ_k of length l_1.

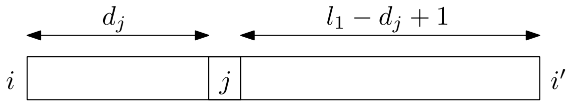

Proof. For each block Θ_k we need to specify a coupling $\Psi_k(x, y)$ of the distributions $P^{[k]}(x, \cdot)$ and $P^{[k]}(y, \cdot)$ for each pair of configurations $(x, y) \in S_i$ and each $i \in V$. Trivially if $i \in \Theta_k$ then the set of H-colourings with positive measure in each distribution is the same and the same H-colouring can be chosen for each distribution. The same holds when i is not on he boundary of Θ_k.

Suppose that i is on the boundary of Θ_k. Let the other site on the boundary of Θ_k be coloured d in both x and y and hence $P^{[k]}(x, \cdot) = D^{(l_1)}_{x_i,d}$ and $P^{[k]}(y, \cdot) = D^{(l_1)}_{y_i,d}$. We then let $\Psi_k(x, y) = \Psi(D^{(l_1)}_{x_i,d}, D^{(l_1)}_{y_i,d})$ which is the coupling constructed in Lemma 11 and gives the stated bounds on the disagreement probabilities. ☐

Remark. It is important to note that, given distinct sites i and i' both on the boundary of Θ_k, we may use a different coupling for $\rho^k_{i,j}$ and $\rho^k_{i',j}$. This is the case since, by definition of ρ, the coupling may depend on both the block and the two initial configurations x and y (which in turn determine i). Since x and y only differ on the colour assigned to site i, the coupling is defined to start from the site in Θ_k immediately adjacent to i, and thus we can use a different coupling for $\rho^k_{i,j}$ and $\rho^k_{i',j}$.

The following technical lemma is required in the proof of Theorem 2.

Lemma 13. *For any* $0 \le p \le 1$ *and* $j, l \in \mathbb{Z}^+$ *where* $l \ge 2j$

$$p^j + p^{l-j+1} \ge p^{j+1} + p^{l-(j+1)+1}.$$

Proof.

$$\begin{aligned} p^j + p^{l-j+1} - p^{j+1} - p^{l-j} &= p^j(1-p) - p^{l-j}(1-p) \\ &= (p^j - p^{l-j})(1-p) \\ &= p^j(1 - p^{l-2j})(1-p) \ge 0 \end{aligned}$$

since $0 \le p \le 1$ where the last equality uses the fact $l \ge 2j$. ☐

We are now ready to prove Theorem 2.

Proof of Theorem 2. We will show that $\alpha < 1$ and then use Theorem 7 to obtain the stated bound on the mixing time. Consider some site $j \in \Theta_k$ and let d_j denote the number of edges between j and the nearest site $i \notin \Theta_k$ on the boundary of Θ_k. Then the distance to the other site, i', on the boundary of Θ_k is $l_1 - d_j + 1$ as shown in Figure 4. Notice that $d_j \le \lceil l_1/2 \rceil$. By Corollary 12 we have

$$\rho^k_{i,j} \le \left(1 - \frac{1}{\Delta_H^2}\right)^{d_j} \text{ and } \rho^k_{i',j} \le \mathbf{1}_{d_j \ge 2}\left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - d_j + 1} + \mathbf{1}_{d_j = 1}\left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - 1}.$$

Now let

$$\alpha_{j,k} = \rho_{i,j}^k + \rho_{i',j}^k \le \left(1 - \frac{1}{\Delta_H^2}\right)^{d_j} + \mathbf{1}_{d_j \ge 2}\left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - d_j + 1} + \mathbf{1}_{d_j = 1}\left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - 1}$$

be the influence on site j. Then

$$\begin{aligned}\alpha &= \max_k \max_{j \in \Theta_k} \alpha_{j,k} \\ &\le \max\left\{ \max_{\lceil \frac{l_1}{2} \rceil \ge d_j \ge 2} \left\{ \left(1 - \frac{1}{\Delta_H^2}\right)^{d_j} + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - d_j + 1} \right\}, \left(1 - \frac{1}{\Delta_H^2}\right) + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - 1} \right\}.\end{aligned}$$

Since $d_j \le \lceil l_1/2 \rceil$ the conditions of Lemma 13 are satisfied for $2 \le d_j \le \lceil l_1/2 \rceil - 1$. In particular taking $d_j = \lceil l_1/2 \rceil - 1$, which satisfies the requirements, gives

$$\left(1 - \frac{1}{\Delta_H^2}\right)^{\lceil l_1/2 \rceil - 1} + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - \lceil l_1/2 \rceil + 2} \ge \left(1 - \frac{1}{\Delta_H^2}\right)^{\lceil l_1/2 \rceil} + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - \lceil l_1/2 \rceil + 1}$$

and hence

$$\begin{aligned}\max_{\lceil \frac{l_1}{2} \rceil \ge d_j \ge 2} \left\{ \left(1 - \frac{1}{\Delta_H^2}\right)^{d_j} + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - d_j + 1} \right\} &\le \left(1 - \frac{1}{\Delta_H^2}\right)^2 + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - 1} \\ &\le \left(1 - \frac{1}{\Delta_H^2}\right) + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - 1}\end{aligned}$$

which gives

$$\begin{aligned}\alpha &\le \left(1 - \frac{1}{\Delta_H^2}\right) + \left(1 - \frac{1}{\Delta_H^2}\right)^{l_1 - 1} \\ &= 1 - \frac{1}{\Delta_H^2} + \left(1 - \frac{1}{\Delta_H^2}\right)^{\lceil \Delta_H^2 \log(\Delta_H^2 + 1) \rceil} \\ &< 1 - \frac{1}{\Delta_H^2} + \frac{1}{\Delta_H^2 + 1} \\ &= 1 - \frac{1}{\Delta_H^2(\Delta_H^2 + 1)}\end{aligned}$$

by substituting the definition of l_1 and using the fact $(1 - 1/x)^x < e^{-1}$ for $x > 0$. The statement of the theorem now follows by Theorem 7. □

We now take a moment to show that we are unable to use Theorem 7 to prove rapid mixing for systematic scan on H-colourings of the n-vertex path for any H that does not have a 2-edge path between all pairs of colours. This motivates the use of path coupling (at the expense of enforcing a specific scan order) in the subsequent section.

Observation 14. *Let $H = (C, E)$ be some fixed and connected graph in which there is no 2-edge path from c_1 to c_2 for some distinct $c_1, c_2 \in C$. Then for any set of m blocks with associated transition matrices $P^{[1]} \dots P^{[m]}$ and any coupling $\Psi_k(x, y)$ for $1 \le k \le m$ and $(x, y) \in S_i$ we have $\alpha \ge 1$ in the unweighted setting.*

Proof. Recall $S_i \subseteq \Omega_\sim^+ \times \Omega_\sim^+$ where $\Omega_\sim^+$ is the set of all configurations (except when H is bipartite in which case $\Omega_\sim^+$ is one of Ω_1^+ and Ω_2^+ as described earlier). Note in particular that any given configuration in $\Omega_\sim^+$ need not be an H-colouring of the n-vertex path. Also recall that $\rho_{i,j}^k$ is the maximum probability of disagreement at j when drawing from a coupling starting from two configurations $(x, y) \in S_i$. Let x be any proper H-colouring with $x_i = c_1$ and y be the configuration with $y_j = x_j$ for $j \neq i$ and $y_i = c_2$ (If H is bipartite then c_2 is from the same colour class of H as c_1). Note that y is not a proper H-colouring as both edges $(y_{i-1}, y_i) \notin E$ and $(y_i, y_{i+1}) \notin E$, otherwise the 2-edge paths $(x_i, x_{i+1} = y_{i+1}, y_i)$ and $(x_i, x_{i-1} = y_{i-1}, y_i)$ would exist in H. However, x and y are both configurations in $\Omega_\sim^+$ and they only differ at the colour of site i so (x, y) is a valid pair in S_i.

Now assume that $\alpha < 1$. Fix some block $\Theta_k = \{i + 1, \ldots, i + l\}$ of length l and let $P^{[k]}$ be the transition matrix associated with Θ_k. Also let $\Psi_k(x, y)$ be any coupling of $P^{[k]}(x, \cdot)$ and $P^{[k]}(y, \cdot)$. Since $\alpha < 1$ it must hold that $\rho_{i,j}^k < 1$ for each $j \in \Theta_k$. In particular $\rho_{i,i+1}^k = \Pr_{(x',y') \in \Psi_k(x,y)}(x'_{i+1} \neq y'_{i+1}) < 1$ and so (letting *adj*(c) denote the set of colours adjacent to c in H) the set *adj*$(c_1) \cap$ *adj*(c_2) must be non-empty since there is a positive probability of assigning the same colour to site $i + 1$ in both distributions. However take any $d \in$ *adj*$(c_1) \cap$ *adj*(c_2), then (c_1, d, c_2) is a 2-edge path from c_1 to c_2 in H contradicting the restriction imposed on H and hence $\alpha \geq 1$. ☐

Remark. It remains to be seen if adding weights will allow a proof in the Dobrushin setting for classes of H not containing 2-edge paths between all colours. However, this can be done using path coupling as we will show in section 3.

3 H-colouring on the path for any H

Recall that $\mathcal{M}_{\text{FixedOrder}}$ is the systematic scan on $\Omega_\sim$ defined as follows. Let $s = 4q + 1$, $\beta = \lceil \log(2sq^s + 1) \rceil q^s$ and $l_2 = 2\beta s$. Then $\mathcal{M}_{\text{FixedOrder}}$ is the systematic scan which performs a heat-bath move on each of the $m_2 + 1 = \lfloor 2n/l_2 \rfloor$ blocks in the order $\Theta_0, \ldots, \Theta_{m_2}$ where

$$\Theta_k = \{k\beta s + 1, \ldots, \min((k + 2)\beta s, n)\}.$$

Note that the size of Θ_{m_2} is at least βs and that every other block is of size l_2. We will prove Theorem 4 which bounds the mixing time of $\mathcal{M}_{\text{FixedOrder}}$. Our method of proof will be path coupling [5] and we begin by establishing some lemmas required to define the coupling we will use in the proof of Theorem 4. The constructions used in the following two lemmas are similar to the ones from Lemma 27 in Dyer et al. [14].

Lemma 15. *If H is not bipartite then for all $c_1, c_2 \in C$ there is an s-edge path in H from c_1 to c_2.*

Proof. Let $c \in C$ be some site on an odd-length cycle in H and let d_1 be the shortest edge-distance from c_1 to c and d_2 the shortest edge-distance from c to c_2. We construct the path as follows. Go from c_1 to c using d_1 edges. If $d_1 + d_2$ is even then go around the cycle using an odd number $q' \leq q$ of edges. Go from c to c_2 in d_2 edges and observe that the constructed path is of odd length. Also the length of the path is at most

$$d_1 + d_2 + q' < 3q.$$

Finally go back and forth on the last edge on the path to make the total length s. $\square$

Lemma 16. *If H is bipartite with colour classes C_1 and C_2 then for all $c_1 \in C_1$ and $c_2 \in C_2$ there is an s-edge path in H from c_1 to c_2.*

Proof. Go from c_1 to c_2 in at most $q-1$ edges and note that the number of edges is odd. Then go back and forth on the last edge to make the total path length equal to s. $\square$

For completeness we present a proof that $\mathcal{M}_{\text{FixedOrder}}$ is ergodic on $\Omega_\sim$.

Lemma 17. *The Markov chain $\mathcal{M}_{\text{FixedOrder}}$ is ergodic on $\Omega_\sim$.*

Proof. Let $P_{\text{FixedOrder}}$ be the transition matrix of $\mathcal{M}_{\text{FixedOrder}}$. We need to show that $\mathcal{M}_{\text{FixedOrder}}$ satisfies the following properties

- irreducible: $P^t_{\text{FixedOrder}}(x, y) > 0$ for each pair $(x, y) \in \Omega_\sim \times \Omega_\sim$ and some integer $t > 0$
- aperiodic: $\gcd\{t : P^t_{\text{FixedOrder}}(x, x) > 0\} = 1$ for each $x \in \Omega_\sim$.

In an application of $P_{\text{FixedOrder}}$ a heat-bath move is made on each block in the order $\Theta_0, \ldots, \Theta_m$. A heat-bath move on any block starting from an H-colouring has a positive probability of self-loop which ensures aperiodicity of the chain. To see that $\mathcal{M}_{\text{FixedOrder}}$ is irreducible consider any pair of H-colourings $(x, y) \in \Omega_\sim \times \Omega_\sim$. We exhibit a sequence of H-colourings $x = \sigma^0, \ldots, \sigma^{m_2+1} = y$ such that $\sigma^k_j = \sigma^{k+1}_j$ for each $0 \le k \le m_2$ and $j \in V \setminus \Theta_k$. Using this sequence we observe that $P_{\text{FixedOrder}}(x, y) > 0$ since, for each $0 \le k \le m_2$, performing a heat-bath move on block Θ_k to $\sigma^k \in \Omega_\sim$ results in the H-colouring $\sigma^{k+1} \in \Omega_\sim$ with positive probability. Recall that $\Theta_k = \{k\beta s + 1, \ldots, \min((k + 2)\beta s, n)\}$. Then let σ^k be given by

$$\sigma^k_i = \begin{cases} y_i & \text{if } 1 \le i \le \min((k+2)\beta s - s + 1, n) \\ x_i & \text{if } (k+2)\beta s + 1 \le i \le n \\ p(i - (k+2)\beta s + s - 1) & \text{if } (k+2)\beta s - s + 1 < i \le \min((k+2)\beta s, n) \end{cases}$$

where $p(j)$ is the j-th in the sequence of colours on the s-edge path in H between $p(0) = y_{(k+2)\beta s - s + 1}$ and $p(s) = x_{(k+2)\beta s + 1}$ given by Lemmas 15 and 16 (since $p(0)$ and $p(s)$ are in opposite colour classes of H in the bipartite case) respectively. $\square$

The following lemma is an analogue of Lemma 13 in Goldberg et al. [19].

Lemma 18. *For any $c_1, c_2, d \in C$ and positive integer $s' \ge s$ such that both $D^{(s')}_{c_1,d}$ and $D^{(s')}_{c_2,d}$ are non-empty there exists a coupling $\psi(D^{(s')}_{c_1,d}, D^{(s')}_{c_2,d})$ of $D^{(s')}_{c_1,d}$ and $D^{(s')}_{c_2,d}$ such that*

$$\Pr_{(x',y') \in \psi(D^{(s')}_{c_1,d}, D^{(s')}_{c_2,d})}(x'_{v_s} \ne y'_{v_s}) \le 1 - \frac{1}{q^s}.$$

Proof. For ease of notation let D_1 denote $D^{(s')}_{c_1,d}$ and D_2 denote $D^{(s')}_{c_2,d}$. For $s' > s$, let n_k be the number of H-colourings on $v_{s+1}, \ldots, v_{s'}$ consistent with v_s being assigned colour $k \in C$ and $v_{s'}$ adjacent to a site (not in L) coloured d. If both $s' = s$ and k is adjacent to d in H then $n_k = 1$. If $s' = s$ but k is not adjacent to d in H then $n_k = 0$. The following definitions are for $i \in \{1,2\}$. Let $l_i(k)$ be the number of H-colourings on $v_1,\dots,v_s$ assigning colour k to site v_s and consistent with v_1 being adjacent to a site (not in L) coloured c_i. We also let Z_i be the set of H-colourings on L with positive measure in D_i and z_i be the size of this set. Note that D_i is the uniform distribution on Z_i so for each $x \in Z_i$ $\mathrm{Pr}_{D_i}(x) = 1/z_i$. For each $k \in C$ let $Z_i(k) \subseteq Z_i$ be the set of H-colourings with positive measure in D_i that assign colour k to site v_s and let $z_i(k)$ be the size of this set. Note that $l_i(k)n_k = z_i(k)$ and $\sum_k z_i(k) = z_i$. Let $C_i^* = \{k \in C \mid z_i(k) > 0\}$ be the set of valid colours for v_s in D_i and let $C^* = C_1^* \cup C_2^*$.

We define a coupling ψ of D_1 and D_2 as follows. Assume without loss of generality that $z_1 \geq z_2$. We create the following mutually exclusive subsets of Z_i. For each $k \in C^*$ let $f(k) = \min(z_1(k), z_2(k))$ and let $F_1(k) = \{\sigma^{(k)}(1),\dots,\sigma^{(k)}(f(k))\} \subseteq Z_1(k)$ be any subset of H-colourings in Z_1 assigning the colour k to site v_s. Also let $F_2(k) = \{\tau^{(k)}(1),\dots,\tau^{(k)}(f(k))\} \subseteq Z_2(k)$ and observe that $F_1(k)$ and $F_2(k)$ are of the same size. We then construct ψ such that for each $k \in C^*$ and $j \in \{1,\dots,f(k)\}$

$$\mathrm{Pr}_{(x',y')\in\psi}(x' = \sigma^{(k)}(j), y' = \tau^{(k)}(j)) = \frac{1}{z_1}.$$

The rest of the coupling is arbitrary. For example let $R_i = Z_i \setminus \left(\bigcup_{k\in C^*} F_i(k)\right)$ be the set of (valid) H-colourings not selected in any of the above subsets of Z_i and the size of R_i be r_i, observing that $r_1 \geq r_2$. Let $R_1' = \{\sigma(1),\dots,\sigma(r_2)\} \subseteq R_1$ and enumerate R_2 such that $R_2 = \{\tau(1),\dots,\tau(r_2)\}$. Then for $1 \leq j \leq r_2$ let

$$\mathrm{Pr}_{(x',y')\in\psi}(x' = \sigma(j), y' = \tau(j)) = \frac{1}{z_1}.$$

Finish off the coupling by, for each pair $(\sigma \in R_1 \setminus R_1', \tau \in Z_2)$ of H-colourings, letting

$$\mathrm{Pr}_{(x',y')\in\psi}(x' = \sigma, y' = \tau) = \frac{1}{z_1 z_2}.$$

From the construction we can verify that the weight of each colouring $x \in Z_1$ in the coupling is $1/z_1$ and the weight of each colouring $y \in Z_2$ is

$$\frac{1}{z_1} + \frac{z_1 - z_2}{z_1 z_2} = \frac{1}{z_2}$$

since the size of $R_1 \setminus R_1'$ is $z_1 - z_2$. This hence completes the construction of the coupling.

We will require the following bounds on $l_i(k)$ for each $k \in C^*$

$$1 \leq l_i(k) \leq q^s. \tag{2}$$

There are at most q colours available for each site in the block and hence at most q^s valid H-colourings of $v_1,\dots,v_s$ which gives the upper bound. We establish the lower bound by showing the existence of an s-edge path in H from both c_1 and c_2 to any $k \in C^*$. Suppose that H is non-bipartite, then Lemma 15 guarantees the existence of an s-edge path in H between any two colours in H, satisfying our requirement.

Now suppose that H is bipartite with colour classes C_1 and C_2. Without loss of generality suppose that $c_1 \in C_1$. Since both D_1 and D_2 are non-empty there exists a

$(2s'+2)$-edge path in H from c_1 to c_2 (via d) so $c_2 \in C_1$. Let $k \in C$ then $k \in C_2$ since there is an s-edge path in H from c_1 to k and s is odd. Lemma 16 implies the existence of an s-edge path between each $c \in C_1$ and each $k \in C_2$ which establishes (2).

Using (2) to see that $n_k \leq f(k) \leq q^s n_k$ for each $k \in C^*$ we have

$$\begin{aligned}
\Pr_{(x',y')\in\psi}(x'_{v_s} = y'_{v_s}) &= \sum_{k\in C^*} \Pr_{(x',y')\in\psi}(x'_{v_s} = y'_{v_s} = k) \\
&\geq \sum_{k\in C^*} \frac{f(k)}{z_1} \\
&\geq \sum_{k\in C^*} \frac{n_k}{\sum_{k'\in C^*} l_1(k')n_k} \\
&\geq \sum_{k\in C^*} \frac{n_k}{q^s \sum_{k'\in C^*} n_{k'}} \\
&= \frac{1}{q^s}
\end{aligned}$$

which completes the proof. □

Lemma 19. *For any $c_1, c_2, d \in C$ and any positive integer $l' \leq l_2$ such that both $D^{(l')}_{c_1,d}$ and $D^{(l')}_{c_2,d}$ are non-empty there exists a coupling Ψ of $D^{(l')}_{c_1,d}$ and $D^{(l')}_{c_2,d}$ in which for $1 \leq j \leq l'$*

$$\Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_j} \neq y'_{v_j}) \leq \left(1 - \frac{1}{q^s}\right)^{\lfloor \frac{j}{s} \rfloor}.$$

Proof. We construct a coupling $\Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})$ of $D^{(l')}_{c_1,d}$ and $D^{(l')}_{c_2,d}$ using the following two step process, based on the recursive coupling in Goldberg et al. [19].

1. If $l' < s$ then couple the distributions any valid way which completes the coupling. Otherwise, couple $D^{(l')}_{c_1,d}(v_s)$ and $D^{(l')}_{c_2,d}(v_s)$ greedily to maximise the probability of assigning the same colour to site v_s in both distributions. Then, independently in each distribution, colour the sites $v_1, \ldots, v_{s-1}$ consistent with the uniform distribution on H-colourings. Note that it is possible to do this since we obtained the colour for site v_s in each distribution from the induced distribution on that site. If $l' = s$ this completes the coupling.

2. If the same colour is assigned to v_s then the remaining sites can be coloured the same way in both distributions since the conditional distributions are the same. Otherwise, for all pairs (c'_1, c'_2) of distinct colours the coupling is completed by recursively constructing a coupling of $\left[D^{(l')}_{c_1,d} \mid v_s = c'_1\right] = D^{(l'-s)}_{c'_1,d}$ and $\left[D^{(l')}_{c_2,d} \mid v_s = c'_2\right] = D^{(l'-s)}_{c'_2,d}$.

This completes the coupling construction and we will prove by strong induction that for $j \in \{1, \ldots, l'\}$

$$\Pr_{(x',y')\in\Psi(D^{(l')}_{c_1,d},D^{(l')}_{c_2,d})}(x'_{v_j} \neq y'_{v_j}) \leq \left(1 - \frac{1}{q^s}\right)^{\lfloor \frac{j}{s} \rfloor}. \tag{3}$$

Firstly the cases $1 \leq j \leq s-1$ are established by observing that $\lfloor j/s \rfloor = 0$ and the probability of disagreement at any site is at most 1. The case $j = s$ is established in

Lemma 18. Now for $s < j \leq l'$, suppose that (3) holds for all positive integers less than j. Let $S_- = \{s, 2s, \dots\}$ and define the quantities j_- and a_j by $j_- = \max\{x \in S_- \mid x < j\} = a_j s$ observing that $1 \leq j - j_- \leq s$. Now

$$\begin{aligned}
&\Pr_{(x',y') \in \Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})}(x'_{v_j} \neq y'_{v_j}) \\
&\quad = \sum_{c'_1, c'_2} \Pr_{(x',y') \in \Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})}(x'_{v_{j_-}} = c'_1, y'_{v_{j_-}} = c'_2) \\
&\qquad \times \Pr_{(x',y') \in \Psi(D^{(l')}_{c_1,d}|v_{j_-} = c'_1, D^{(l')}_{c_2,d}|v_{j_-} = c'_2)}(x'_{v_j} \neq y'_{v_j}) \\
&\quad = \sum_{c'_1, c'_2} \Pr_{(x',y') \in \Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})}(x'_{v_{j_-}} = c'_1, y'_{v_{j_-}} = c'_2) \\
&\qquad \times \Pr_{(x',y') \in \Psi(D^{(l'-j_-)}_{c'_1,d}, D^{(l'-j_-)}_{c'_2,d})}(x'_{v_{j-j_-}} \neq y'_{v_{j-j_-}}).
\end{aligned}$$

Observe that for any pair (c'_1, c'_2) of colours, if the probabilities of assigning c'_1 to v_{j_-} in $D^{(l')}_{c_1,d}$ and c'_2 to v_{j_-} in $D^{(l')}_{c_2,d}$ are both non-zero then the distributions $D^{(l'-j_-)}_{c'_1,d}$ and $D^{(l'-j_-)}_{c'_2,d}$ are both non-empty and hence, using Lemma 18 for $l' - j_- \geq s$ and upper-bounding probability of disagreement by one otherwise, we get

$$\begin{aligned}
&\Pr_{(x',y') \in \Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})}(x'_{v_j} \neq y'_{v_j}) \\
&\quad \leq \sum_{c'_1, c'_2} \Pr_{(x',y') \in \Psi(D^{(l')}_{c_1,d}, D^{(l')}_{c_2,d})}(x'_{v_{j_-}} = c'_1, y'_{v_{j_-}} = c'_2) \left(\mathbf{1}_{j-j_-=s}(1 - 1/q^s) + \mathbf{1}_{j-j_- \neq s}\right) \\
&\quad \leq \begin{cases} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{j_-}{s} \right\rfloor + 1} & \text{if } j - j_- = s \\ \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{j_-}{s} \right\rfloor} & \text{if } j - j_- \neq s \end{cases}
\end{aligned} \tag{4}$$

where last inequality is the inductive hypothesis since $j_- < j$.

First consider the case $j - j_- \neq s$ in which we have $j_- + b = j$ for some $1 \leq b \leq s - 1$. Then

$$\left\lfloor \frac{j_- - 1}{s} \right\rfloor = \left\lfloor \frac{a_j s - 1}{s} \right\rfloor = a_j - 1 < a_j = \left\lfloor \frac{a_j s}{s} \right\rfloor = \left\lfloor \frac{j_-}{s} \right\rfloor$$

and so for $1 \leq b \leq s - 1$

$$\left\lfloor \frac{j_- + b}{s} \right\rfloor = \left\lfloor \frac{j_-}{s} \right\rfloor$$

which implies that

$$\left\lfloor \frac{j_-}{s} \right\rfloor = \left\lfloor \frac{j}{s} \right\rfloor. \tag{5}$$

Now suppose $j - j_- = s$ which substituting for j_- gives

$$\left\lfloor \frac{j_-}{s} \right\rfloor = \left\lfloor \frac{j - s}{s} \right\rfloor = \left\lfloor \frac{j}{s} \right\rfloor - 1. \tag{6}$$

Substituting (5) and (6) in (4) completes the proof. □

We are now ready to define the coupling of the distributions of configurations obtained from one complete scan of the Markov chain $\mathcal{M}_{\text{FixedOrder}}$. The coupling is defined for pairs $(x, y) \in S_i$. We will let (x', y') denote the pair of configurations after one complete scan of $\mathcal{M}_{\text{FixedOrder}}$ starting from (x, y) and let (x^k, y^k) be the pair of configurations obtained by updating blocks $\Theta_0, \ldots, \Theta_{k-1}$ starting from $(x, y) = (x^0, y^0)$. Observe that (x', y') is obtained by updating block Θ_{m_2} from the pair (x^{m_2}, y^{m_2}).

The coupling for updating block Θ_k is defined as follows. Let i and i' be the sites on the boundary of Θ_k. The order of the scan will ensure that at most one of the boundaries is a disagreement in (x^k, y^k), so we only need to define the coupling for boundaries disagreeing on at most one end of Θ_k; suppose without loss of generality that $x^k_{i'} = y^k_{i'} = d$ for some $d \in C$. Firstly, if $x^k_i = y^k_i$ then the set of valid configurations arising from updating Θ_k is the same in both distributions and we use the identity coupling.

Otherwise $x^k_i \neq y^k_i$. If H is not bipartite then Lemma 15 implies the existence of a $(m_2 + 1)$-edge path between both x^k_i and d and between y^k_i and d. If H is bipartite then x^k_i and y^k_i are in the same colour class but d is in the opposite colour class of H since l_2 is even. Lemma 16 implies the existence of a $(m_2 + 1)$-edge path between both x^k_i and d and between y^k_i and d. Hence both distributions $D^{(l_2)}_{x^k_i,d}$ and $D^{(l_2)}_{y^k_i,d}$ are non-empty and we obtain (x^{k+1}, y^{k+1}) from $\Psi(D^{(l_2)}_{x^k_i,d}, D^{(l_2)}_{y^k_i,d})$ which is the coupling constructed in Lemma 19. Note that if $k = m_2$ (i.e. the block is the last block which may not be of size l_2) then both distributions remain (trivially) non-empty. For ease of reference we state the following corollary of Lemma 19.

Corollary 20. *For any two sites $v, u \in V$ let $d(v, u)$ denote the edge distance between them. For any block Θ_k let i and i' be the sites on the boundary of Θ_k and suppose that $x^k_{i'} = y^k_{i'} = d$ for any $d \in C$. Obtain (x^{k+1}, y^{k+1}) from the above coupling. Then for any $j \in \Theta_k$*

$$\Pr(x^{k+1}_j \neq y^{k+1}_j) \leq \begin{cases} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{d(i,j)}{s} \right\rfloor} & \text{if } x^k_i \neq y^k_i \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 21. *For any positive integers s, k, x*

$$\sum_{j=1}^{sk} \left(1 - \frac{1}{x}\right)^{\left\lfloor \frac{j}{s} \right\rfloor} < sx.$$

Proof.

$$\sum_{j=1}^{sk} \left(1 - \frac{1}{x}\right)^{\left\lfloor \frac{j}{s} \right\rfloor} = (s-1) + s\sum_{j=1}^{k-1} \left(1 - \frac{1}{x}\right)^j + \left(1 - \frac{1}{x}\right)^k < s\sum_{j\geq 0} \left(1 - \frac{1}{x}\right)^j < sx.$$

□

The following lemma implies Theorem 4 by Theorem 9 (path coupling).

Lemma 22. *Suppose that $(x, y) \in S_i$ and obtain (x', y') by one complete scan of $\mathcal{M}_{\text{FixedOrder}}$. Then*

$$\mathbf{E}\left[\text{Ham}(x', y')\right] < 1 - \frac{1}{4sq^s + 2}.$$

Figure 5: Site i is on the boundary of Θ_a and is not contained in any block $\Theta_{a'}$ with $a' < a$.

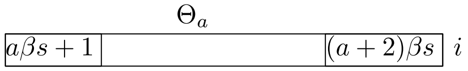

Proof. First suppose that i is not on the boundary of any block and that Θ_b is the first block containing i. In this case Corollary 20 gives us $\Pr(x_i^{b+1} \neq y_i^{b+1}) = 0$ and so $\mathrm{Ham}(x', y') = 0$.

Now suppose that i is on the boundary of some block Θ_a. Recall the definition of a block

$$\Theta_k = \{k\beta s + 1, \ldots, \min(k\beta s + 2\beta s, n)\}.$$

If i is also contained in a block $\Theta_{a'}$ with $a' < a$ then Corollary 20 gives $\Pr(x_i^{a'+1} \neq y_i^{a'+1}) = 0$ and hence $\mathrm{Ham}(x', y') = 0$.

If site i is not updated before Θ_a then $i = (a+2)\beta s + 1$ as shown in Figure 5 and the disagreement percolates through the sites in Θ_a during the update of Θ_a. Using Corollary 20 we have for $j \in \Theta_a$

$$\Pr(x_j^{a+1} \neq y_j^{a+1}) \leq \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{i-j}{s} \right\rfloor} \tag{7}$$

in particular, the sites in $\Theta_a \setminus \Theta_{a+1} = \{a\beta s + 1, \ldots (a+1)\beta s\}$ will not get updated again during the scan and hence for $j \in \Theta_a \setminus \Theta_{a+1}$

$$\Pr(x'_j \neq y'_j) \leq \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{(a+2)\beta s+1-j}{s} \right\rfloor}. \tag{8}$$

Now consider the update of any block Θ_k from the pair of configurations (x^k, y^k) where $k > a$. There cannot be a disagreement at site $(k+2)\beta s + 1$ since that site has not been updated (and it was not the initial disagreement) so the only site on the boundary of Θ_k that could be a disagreement in (x^k, y^k) is $k\beta s$. Hence from Corollary 20, for $j \in \{k\beta s + 1, \ldots, \min((k+2)\beta s, n)\}$

$$\Pr(x_j^{k+1} \neq y_j^{k+1} \mid x_{k\beta s}^k \neq y_{k\beta s}^k) \leq \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{j-k\beta s}{s} \right\rfloor}. \tag{9}$$

We show by induction on k that for $a + 1 \leq k \leq m_2$

$$\Pr(x_{k\beta s}^k \neq y_{k\beta s}^k) \leq \left(1 - \frac{1}{q^s}\right)^{\beta(k-a)}. \tag{10}$$

The base case, $k = a + 1$ follows from (7) since $j = k\beta s = (a+1)\beta s = a\beta s + \beta s \in \Theta_a$.

Now suppose that (10) is true for $k-1$. Then

$$\begin{aligned}
\Pr(x^k_{k\beta s} \neq y^k_{k\beta s}) &= \Pr(x^k_{k\beta s} \neq y^k_{k\beta s} \mid x^{k-1}_{(k-1)\beta s} \neq y^{k-1}_{(k-1)\beta s}) \Pr(x^{k-1}_{(k-1)\beta s} \neq y^{k-1}_{(k-1)\beta s}) \\
&\leq \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{k\beta s - (k-1)\beta s}{s} \right\rfloor} \left(1 - \frac{1}{q^s}\right)^{\beta(k-a-1)} \\
&= \left(1 - \frac{1}{q^s}\right)^{\beta} \left(1 - \frac{1}{q^s}\right)^{\beta(k-a-1)} \\
&= \left(1 - \frac{1}{q^s}\right)^{\beta(k-a)}
\end{aligned}$$

using the inductive hypothesis and (9).

Now for each site $j \geq (a+1)\beta s + 1$, that is site j is updated at least once following block Θ_a, write $j = k_j \beta s + b_j$ with $1 \leq b_j \leq \beta s$ where k_j denotes is the index of the block in which j is last updated.

$$\begin{aligned}
\Pr(x'_j \neq y'_j) &= \Pr(x^{k_j+1}_j \neq y^{k_j+1}_j) \\
&\leq \Pr(x^{k_j+1}_j \neq y^{k_j+1}_j \mid x^{k_j}_{\beta k_j s} \neq y^{k_j}_{\beta k_j s}) \Pr(x^{k_j}_{\beta k_j s} \neq y^{k_j}_{\beta k_j s}).
\end{aligned}$$

We can then apply (9) to the first component of the product since $j \in \{k_j\beta s + 1, \ldots, \min(k_j \beta s + 2\beta s, n)\}$ and (10) to the second since $a + 1 \leq k_j \leq m_2$ to get

$$\Pr(x'_j \neq y'_j) \leq \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{b_j}{s} \right\rfloor} \left(1 - \frac{1}{q^s}\right)^{\beta(k_j - a)}.$$

Then, using linearity of expectation and (8), we have

$$\begin{aligned}
\mathbf{E}\left[\mathrm{Ham}(x', y')\right] &= \sum_j \Pr(x'_j \neq y'_j) \\
&= \sum_{j \in \Theta_a \setminus \Theta_{a+1}} \Pr(x'_j \neq y'_j) + \sum_{j \in \bigcup_{k \geq a+1} \Theta_k} \Pr(x'_j \neq y'_j) \\
&\leq \sum_{j = as\beta + 1}^{(a+1)\beta s} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{(a+2)\beta s + 1 - j}{s} \right\rfloor} + \sum_{k_j = a+1}^{m_2} \sum_{b_j = 1}^{\beta s} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{b_j}{s} \right\rfloor} \left(1 - \frac{1}{q^s}\right)^{\beta(k_j - a)} \\
&= \sum_{r=1}^{\beta s} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{\beta s + r}{s} \right\rfloor} + \sum_{k_j = a+1}^{m_2} \left(1 - \frac{1}{q^s}\right)^{\beta(k_j - a)} \sum_{b_j = 1}^{\beta s} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{b_j}{s} \right\rfloor} \\
&< \left(1 - \frac{1}{q^s}\right)^{\beta} \sum_{r=1}^{\beta s} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{r}{s} \right\rfloor} + \sum_{t \geq 1} \left(\left(1 - \frac{1}{q^s}\right)^{\beta}\right)^{t} \sum_{b_j = 1}^{\beta s} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{b_j}{s} \right\rfloor} \\
&< \left(1 - \frac{1}{q^s}\right)^{\beta} s q^s + \frac{\left(1 - \frac{1}{q^s}\right)^{\beta} s q^s}{1 - \left(1 - \frac{1}{q^s}\right)^{\beta}}
\end{aligned}$$

where the last inequality uses Lemma 21 and the sum of a geometric progression. Substituting the definition of β and using the fact $(1-1/x)^x < e^{-1}$ for $x > 0$ we get

$$\begin{aligned}
\mathbf{E}\left[\mathrm{Ham}(x',y')\right] &< \left(1-\frac{1}{q^s}\right)^{\lceil\log(2sq^s+1)\rceil q^s} sq^s + \frac{\left(1-\frac{1}{q^s}\right)^{\lceil\log(2sq^s+1)\rceil q^s} sq^s}{1-\left(1-\frac{1}{q^s}\right)^{\lceil\log(2sq^s+1)\rceil q^s}} \\
&< \frac{sq^s}{e^{\lceil\log(2sq^s+1)\rceil}} + \frac{sq^s}{e^{\lceil\log(2sq^s+1)\rceil}\left(1-e^{-\lceil\log(2sq^s+1)\rceil}\right)} \\
&= \frac{sq^s}{e^{\lceil\log(2sq^s+1)\rceil}} + \frac{sq^s}{e^{\lceil\log(2sq^s+1)\rceil}-1} \\
&\leq \frac{sq^s}{2sq^s+1} + \frac{sq^s}{2sq^s} \\
&= 1-\frac{1}{4sq^s+2}
\end{aligned}$$

which completes the proof. □

4 H-colouring using a random update Markov chain

Recall that the random update Markov chain $\mathcal{M}_{\mathrm{RND}}$ on $\Omega_\sim$ is defined as follows. We again let $s = 4q+1$ and we define $\gamma = 2q^s+1$. We then define a set of $n+s\gamma-1$ blocks of size at most $s\gamma$ as follows.

$$\Theta_k = \begin{cases} \{k,\ldots,\min(k+s\gamma-1,n)\} & \text{when } k\in\{1,\ldots,n\} \\ \{1,\ldots,n+s\gamma-k\} & \text{when } k\in\{n+1,\ldots,n+s\gamma-1\} \end{cases}$$

By construction of the set of blocks each site is adjacent to *at most* two blocks and furthermore each site is contained in *exactly* $s\gamma$ blocks. One step of $\mathcal{M}_{\mathrm{RND}}$ consists of selecting a block uniformly at random and performing a heat-bath update on it. We will prove (using path coupling) Theorem 6 namely that $\mathcal{M}_{\mathrm{RND}}$ mixes in $O(n\log n)$ updates for any H.

We begin by defining the required coupling. For a pair of configurations $(x,y)\in S_i$ we obtain the pair (x',y') by one step of $\mathcal{M}_{\mathrm{RND}}$. That is we select a block uniformly at random and perform a heat bath move on that block. We can again use Lemma 19 from Section 3 to construct the required coupling for updating block Θ_k since the definition of s is the same in both Markov chains. If i is not on the boundary of Θ_k then the sets of valid H-colourings of Θ_k are the same in both distributions and we use the identity coupling. If i is on the boundary of Θ_k then we let the other site on the boundary be coloured d in both x and y. We then obtain (x',y') from $\Psi(D^{(s\gamma)}_{x_i,d}, D^{(s\gamma)}_{y_i,d})$ which is the coupling constructed in Lemma 19. The disagreement probabilities are summarised in the following corollary (of Lemma 19).

Corollary 23. *For any two sites $v,u\in V$ let $d(v,u)$ denote the edge distance between them. Suppose that a block Θ_k has been selected to be updated. For any pair $(x,y)\in S_i$*

obtain (x', y') from the above coupling. Then for any $j \in \Theta_k$

$$\Pr(x'_j \neq y'_j) \leq \begin{cases} \left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{d(i,j)}{s} \right\rfloor} & \text{if } i \text{ is on the boundary of } \Theta_k \\ 0 & \text{otherwise.} \end{cases}$$

The following lemma implies Theorem 6 by Theorem 9 (path coupling).

Lemma 24. *Suppose that $(x, y) \in S_i$ and obtain (x', y') by one step of $\mathcal{M}_{\mathrm{RND}}$. Then*

$$\mathbf{E}\left[\mathrm{Ham}(x', y')\right] < 1 - \frac{s}{n + 2sq^s + s - 1}.$$

Proof. There are $s\gamma$ blocks containing site i and if such a block is selected then $\mathrm{Ham}(x', y') = 0$. There are at most 2 blocks adjacent to site i and if such a block is selected then the discrepancy percolates in the block according to the probabilities stated in Corollary 23. This leaves $n + s\gamma - 1 - s\gamma - 2 = n - 3$ blocks that leave the Hamming distance unchanged. Hence, using Lemma 21, we have

$$\begin{aligned} \mathbf{E}\left[\mathrm{Ham}(x', y')\right] &\leq \frac{2}{n + s\gamma - 1}\left(1 + \sum_{j=1}^{\gamma s}\left(1 - \frac{1}{q^s}\right)^{\left\lfloor \frac{j}{s} \right\rfloor}\right) + \frac{n-3}{n + s\gamma - 1} \\ &< \frac{n-1}{n + s\gamma - 1} + \frac{2sq^s}{n + s\gamma - 1} \\ &= \frac{2sq^s + n - 1}{2sq^s + n - 1 + s} = 1 - \frac{s}{2sq^s + n - 1 + s} \end{aligned}$$

by substituting the definition of γ. □

Acknowledgments

I am grateful to Leslie Goldberg for several useful discussions regarding technical issues and for providing detailed and helpful comments on a draft of this article.

References

[1] Dimitris Achlioptas, Mike Molloy, Cristopher Moore, and Frank Van Bussel. Sampling grid colourings with fewer colours. In *Proc. of the 6th Latin American Symposium on Theoretical Informatics (LATIN'04)*, pages 80–89, Buenos Aires, Argentina, 2004.

[2] David J Aldous. Random walks on finite groups and rapidly mixing markov chains. In *Séminaire de probabilités XVII*, pages 243–297. Springer-Verlag, 1983.

[3] Yali Amit. Convergence properties of the Gibbs sampler for pertubations of gaussians. *The Annals of Statistics*, 24(1):122–140, 1996.

[4] Magnus Bordewich, Martin Dyer, and Marek Karpinski. Stopping times, metrics and approximate counting. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP*, volume 4051 of *Lecture Notes in Computer Science*, pages 108–119. Springer, 2006.

[5] Russ Bubley and Martin Dyer. Path coupling: a technique for proving rapid mixing in Markov chains. In *38th Annual Symposium on Foundations of Computer Science*, pages 223–231, 1997.

[6] Robert Burton and Jeffrey Steif. Nonuniqueness of measures of maximal entropy for subshifts of finite type. *Ergodic Theory and Dynamical Systems*, 14(2):213–236, 1994.

[7] Colin Cooper, Martin Dyer, and Alan Frieze. On Markov chains for randomly H-colouring a graph. *Journal of Algorithms*, 39(1):117–134, 2001.

[8] Mary Kathryn Cowles and Bradlet P. Carlin. Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of The American Statistical Association*, 91(434):883–904, 1996.

[9] Persi Diaconis and Arun Ram. Analysis of systematic scan Metropolis algorithms using Iwahoti-Hecke algebra techniques. *Michigan Mathematical Journal*, 48:157–190, 2000.

[10] Martin Dyer, Alan Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM Journal Computing*, 31(5):1527–1541, 2002.

[11] Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, and Mark Jerrum. On the relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2003.

[12] Martin Dyer, Leslie Ann Goldberg, and Mark Jerrum. Counting and sampling H-colourings. *Information and Computation*, 189:1–16, 2004.

[13] Martin Dyer, Leslie Ann Goldberg, and Mark Jerrum. Dobrushin conditions and systematic scan. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick, editors, *APPROX-RANDOM*, volume 4110 of *Lecture Notes in Computer Science*, pages 327–338. Springer, 2006.

[14] Martin Dyer, Leslie Ann Goldberg, and Mark Jerrum. Systematic scan and sampling colourings. *Annals of Applied Probability*, 16(1):185–230, 2006.

[15] Martin Dyer and Catherine S. Greenhill. The complexity of counting graph homomorphisms. *Random Structures and Algorithms*, 17:260–289, 2000.

[16] Martin Dyer and Catherine S. Greenhill. On Markov chains for independent sets. *J. Algorithms*, 35(1):17–49, 2000.

[17] Anna Galluccio, Pavol Hell, and Jaroslav Nešetřil. The complexity of H-colouring of bounded degree graphs. *Discrete Mathematics*, 222:101–109, 2000.

[18] Leslie Ann Goldberg, Steven Kelk, and Mike Paterson. The complexity of choosing an H-colouring (nearly) uniformly at random. *SICOMP*, 33(2):416–432, 2004.

[19] Leslie Ann Goldberg, Russ Martin, and Mike Paterson. Strong spatial mixing for lattice graphs with fewer colours. *SICOMP*, 35(2):486–517, 2005.

[20] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-colouring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990.

[21] Mark Jerrum. A very simple algorithm for estimating the number of k-colourings of a low-degree graph. *Random Structures and Algorithms*, 1995.

[22] Claire Kenyon, Elchanan Mossel, and Yuval Peres. Glauber dynamics on trees and hyperbolic graphs. In *Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 568–578, 2001.

[23] Michael Luby and Eric Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets: Part I. *Random Structures and Algorithms*, 15(3–4):229–241, 1999.

[24] Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Glauber dynamics on trees: Boundary conditions and mixing time. *Communications in Mathematical Physics*, 250(2):301–334, 2004.

[25] Kasper Pedersen. Dobrushin conditions for systematic scan with block dynamics (extended abstract). To appear in MFCS, 2007.

[26] Jesus Salas and Alan D Sokal. Absence of phase transition for antiferromagnetic potts models via the dobrushin uniqueness theorem. *Journal of Statistical Physics*, pages 551–579, 1997.

[27] Eric Vigoda. Improved bounds for sampling colourings. *J. Math. Phys*, 2000.

[28] Dror Weitz. Combinatorial criteria for uniqueness of Gibbs measures. *Random Structures and Algorithms*, 27(4):445–475, 2005.

[29] Dror Weitz. Counting independent sets up to the tree threshold. In *STOC*, pages 140–149, 2006.

[30] Benjamin Widom and John S. Rowlinson. New model for the study of liquid-vapour phase transition. *The Journal of Chemical Physics*, 52(4):1670–1684, 1970.